
Variable sigma Gaussian processes: An expectation propagation perspective

Yuan (Alan) Qi
CS & Statistics Departments, Purdue University
alanqi@cs.purdue.edu

Ahmed H. Abdel-Gawad
ECE Department, Purdue University
aabdelga@purdue.edu

Thomas P. Minka
Microsoft Research Cambridge, UK
minka@microsoft.com

Abstract

Gaussian processes (GPs) provide a probabilistic nonparametric representation of functions in regression, classification, and other problems. Unfortunately, exact learning with GPs is intractable for large datasets. A variety of approximate GP methods have been proposed that essentially map the large dataset into a small set of basis points. The most advanced of these, the variable-sigma GP (VSGP) (Walder et al., 2008), allows each basis point to have its own length scale. However, VSGP was only derived for regression. We describe how VSGP can be applied to classification and other problems, by deriving it as an expectation propagation algorithm. In this view, sparse GP approximations correspond to a KL-projection of the true posterior onto a compact exponential family of GPs. VSGP constitutes one such family, and we show how to enlarge this family to get additional accuracy. In particular, we show that endowing each basis point with its own *full covariance matrix* provides a significant increase in approximation power.

1 Introduction

Gaussian processes (GP) are powerful nonparametric Bayesian approach to modelling unknown functions. As such, they can be directly used for classification and regression (Rasmussen & Williams, 2006), or embedded into a larger model such as factor analysis (Teh et al., 2005), relational learning (Chu et al., 2006), or reinforcement learning (Deisenroth et al., 2009). Unfortunately, the cost of GPs can be prohibitive for large datasets. Even for the regression case where the GP prediction formula is analytic, training the exact GP model with N points demands an $O(N^3)$ cost for inverting the covariance matrix and predicting a new output requires $O(N^2)$ cost in addition to storing all of the training points.

Ideally, we would like a compact representation, much smaller than the number of training points, of the posterior distribution for the unknown function. This compact representation could be used to summarize training data for Bayesian learning, or it could be passed around as a message in order to do inference in a larger probabilistic model. One successful approach is to map the training data into a small set of basis points, then compute the exact posterior that results from those points. The basis points could literally be a subset of the training instances (Csato, 2002; Lawrence et al., 2002) or they could be artificial “pseudo-inputs” that represent the training set (Snelson & Ghahramani, 2006). Furthermore, the GP that is used to for inference on the basis points need not match the original GP. In particular, as shown by Walder et al. (2008), the GP applied to the basis points should have a longer length scale than the original (since the data is now sparser). Their “variable sigma Gaussian process” (VSGP) algorithm (Walder et al., 2008) allows a different length scale for each basis point.

For classification problems, a further approximation step is needed. The training set needs to not only be reduced but also transformed from classification data to regression data. In this paper, we show how both of these approximation steps can be accomplished in the same framework of expectation propagation (EP). The novel contributions of this paper include:

- Derive VSGP and other sparse GP algorithms as EP algorithms employing a particular choice of approximating family,
- Extend VSGP to use a full covariance matrix around each basis point,
- Show how to apply the VSGP approximation to classification and other likelihoods, and
- Evaluate the benefit of full covariance matrices on classification and regression problems (section 5).

2 Gaussian process regression and classification

We denote N independent and identically distributed samples as $\mathcal{D} = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\}_N$, where \mathbf{x}_i is a d dimensional input and y_i is a scalar output. We assume there is a latent function f that we are modeling and the noisy realization of latent function f at \mathbf{x}_i is y_i .

A Gaussian process places a prior distribution over the latent function f . Its projection at the samples $\{\mathbf{x}_i\}$ defines a joint Gaussian distribution:

$$p(\mathbf{f}) = \mathcal{N}(\mathbf{f} | \mathbf{m}^0, K)$$

where $m_i^0 = m^0(\mathbf{x}_i)$ is the mean function and $K_{ij} = K(\mathbf{x}_i, \mathbf{x}_j)$ is the covariance function, which encodes the prior notation of smoothness. Normally the mean function is simply set to be zero and we follow this tradition in this paper. A typical kernel covariance function is the squared exponential, also know as Radial Basis Function (RBF),

$$k(\mathbf{x}, \mathbf{x}') = \exp\left(-\frac{\|\mathbf{x}' - \mathbf{x}\|^2}{2\sigma^2}\right), \quad (1)$$

where σ is a hyperparameter.

For regression, we use a Gaussian likelihood function

$$p(y_i | f) = \mathcal{N}(y_i | f(\mathbf{x}_i), v_y) \quad (2)$$

where v_y is the observation noise. For classification, the data likelihood has the form

$$p(y_i | f) = (1 - \epsilon)\phi(f(\mathbf{x}_i)y_i) + \epsilon\phi(-f(\mathbf{x}_i)y_i) \quad (3)$$

where ϵ models the labeling error and $\phi(\cdot)$ is a nonlinear function, ie., a cumulative Gaussian distribution or a step function such that $\phi(f(x_i)) = 1$ if $f(x_i) \leq 0$ and $\phi(f(x_i)) = 0$ otherwise.

Given the Gaussian process prior over f and the data likelihood, the posterior process is

$$p(f | \mathcal{D}, \mathbf{t}) \propto GP(f | 0, K) \prod_{i=1}^N p(y_i | f) \quad (4)$$

Since the Gaussian process is grounded on the N examples, they are called the basis points.

For the regression problem, the posterior process has an analytical form. But to make a prediction on a new sample, we need to invert a N by N matrix. If the training set is big, this matrix inversion will be too costly. For classification or other nonlinear problems, the computational cost is even higher since we do not have a analytical solution to the posterior process and the complexity of the process grows with the number of training samples.

3 Adaptive sparse Gaussian processes

In this section we present the novel sparse Gaussian process model coupled with an efficient inference algorithm.

3.1 Model

To save the computational and memory cost, we approximate the exact Gaussian process model (4) by a sparse model parameterized by $(\mathbf{u}, \mathbf{b}, \mathbf{c}, \mathbf{\Lambda})$:

$$q(f) \propto GP(f|0, K) \prod_{k=1}^M \mathcal{N}(u_k | \int f(\mathbf{x}) \mathcal{N}(\mathbf{x}|b_k, c_k) d\mathbf{x}, \lambda_k^{-1}) \quad (5)$$

In this sparse model, we can change the number of basis points, M , to regulate the model complexity. Normally, we set $M \ll N$. Importantly, the basis point b_k is blurred by the Gaussian distribution $\mathcal{N}(\mathbf{x}|b_k, c_k)$, which can represent how the data distributed *locally* around b_k . The parameter u_k represents a virtual regression target for b_k , and λ_k is its precision.

The distribution (5) is a Gaussian process, so it has a mean and covariance function. To determine these, first consider the following functional:

$$Z(m^0) = \int GP(f|m^0, K) \mathcal{N}(\mathbf{u} | \mathbf{g}_B(f), \mathbf{\Lambda}^{-1}) df \quad (6)$$

where $\mathbf{u} = (u_1, \dots, u_M)$, the subscript B denote the basis set (b_1, \dots, b_M) , and

$$\mathbf{g}_B(f) = \left[\int f(\mathbf{x}) \mathcal{N}(\mathbf{x}|b_1, c_1) d\mathbf{x}, \dots, \int f(\mathbf{x}) \mathcal{N}(\mathbf{x}|b_M, c_M) d\mathbf{x} \right]^T \quad (7)$$

$$\mathbf{\Lambda}^{-1} = \text{diag}(\lambda_1^{-1}, \lambda_2^{-1}, \dots, \lambda_M^{-1}) \quad (8)$$

Now $\mathbf{g}_B(f)$ follows a joint Gaussian distribution with the mean $\tilde{\mathbf{m}}_B^0$ and $\tilde{\mathbf{V}}_B^0$:

$$\tilde{m}_{Bj}^0 = \int \int f(\mathbf{x}) \mathcal{N}(\mathbf{x}|b_j, c_j) GP(f|m^0, K) d\mathbf{x} df = \int m^0(\mathbf{x}) \mathcal{N}(\mathbf{x}|b_j, c_j) d\mathbf{x} \quad (9)$$

$$\tilde{V}_{Bij}^0 = \int \int \mathcal{N}(\mathbf{x}|b_i, c_i) K(\mathbf{x}, \mathbf{x}') \mathcal{N}(\mathbf{x}'|b_j, c_j) d\mathbf{x} d\mathbf{x}' \quad (10)$$

Therefore, we can compute $Z(m^0)$ as follows:

$$Z(m^0) = \int \mathcal{N}(\mathbf{u} | \mathbf{g}_B, \mathbf{\Lambda}^{-1}) \mathcal{N}(\mathbf{g}_B | \tilde{\mathbf{m}}_B^0, \tilde{\mathbf{V}}_B^0) df = \mathcal{N}(\mathbf{u} | \tilde{\mathbf{m}}_B^0, \beta^{-1}) \quad (11)$$

where $\beta = (\tilde{\mathbf{V}}_B^0 + \mathbf{\Lambda}^{-1})^{-1}$. Define $\alpha = \beta \mathbf{u}$. Then the mean and the covariance functions of the sparse model are characterized by α and β . In particular, we have the following theorem to describe the relationship between them.

Theorem 1 *The posterior process $q(f)$ defined in (5) has the mean function $m(\mathbf{x})$ and covariance function $V(\mathbf{x}, \mathbf{x}')$:*

$$m(\mathbf{x}) = \tilde{K}(\mathbf{x}, B) \alpha \quad (12)$$

$$V(\mathbf{x}, \mathbf{x}') = K(\mathbf{x}, \mathbf{x}') - \tilde{K}(\mathbf{x}, B) \beta \tilde{K}(B, \mathbf{x}') \quad (13)$$

where $\tilde{K}(\mathbf{x}, B) = [\tilde{K}(\mathbf{x}, b_1), \dots, \tilde{K}(\mathbf{x}, b_j), \dots, \tilde{K}(\mathbf{x}, b_M)]$, $\tilde{K}(B, \mathbf{x}) = (\tilde{K}(\mathbf{x}, B))^T$, and $\tilde{K}(\mathbf{x}, b_j) = \int K(\mathbf{x}, \mathbf{x}') \mathcal{N}(\mathbf{x}'|b_j, c_j) d\mathbf{x}'$.

When using the RBF covariance function (1), we have

$$\tilde{K}(x, B) = (2\pi\sigma^2)^{M/2} [\mathcal{N}(\mathbf{x}|b_1, c_1 + \sigma^2 \mathbf{I}), \dots, \mathcal{N}(\mathbf{x}|b_M, c_M + \sigma^2 \mathbf{I})]. \quad (14)$$

Proof: First, consider the minimization of the KL divergence between the posterior process $q(f) \propto GP(f|m^0, K) \mathcal{N}(\mathbf{u} | \mathbf{g}_B(f), \mathbf{\Lambda}^{-1})$ for the sparse model and a process $\bar{q}(f)$ in the exponential family. Since $q(f)$ also belongs to the exponential family, this minimization will achieve the optimal solution, i.e., $\bar{q}(f) = q(f)$.

Now to obtain the mean function $m(\mathbf{x})$ and the covariance function $V(\mathbf{x}, \mathbf{x}')$ for $q(f)$, we can solve $\bar{q}(f)$ by the KL minimization. This leads to the following moment matching equations:

$$m(\mathbf{x}) = m^0(\mathbf{x}) + \int K(\mathbf{x}, \mathbf{a}') \frac{d \log Z}{dm^0(\mathbf{a}')} d\mathbf{a}' \quad (15)$$

$$V(\mathbf{x}, \mathbf{x}') = K(\mathbf{x}, \mathbf{x}') + \int \int K(\mathbf{x}, \mathbf{a}) \frac{d^2 \log Z}{dm^0(\mathbf{a})dm^0(\mathbf{a}')} K(\mathbf{a}', \mathbf{x}') d\mathbf{a} d\mathbf{a}' \quad (16)$$

Based on (11), it is easy to obtain

$$\frac{d \log Z}{dm^0(\mathbf{a})} = \frac{d}{d\tilde{\mathbf{m}}_B} \left(-\frac{1}{2}(\mathbf{u} - \tilde{\mathbf{m}}_B)^T \beta (\mathbf{u} - \tilde{\mathbf{m}}_B) \right) \frac{d\tilde{\mathbf{m}}_B}{dm^0(\mathbf{a})} \quad (17)$$

$$= [\mathcal{N}(\mathbf{x}|b_1, c_1), \dots, \mathcal{N}(\mathbf{x}|b_M, c_M)] \beta (\mathbf{u} - \tilde{\mathbf{m}}_B) \quad (18)$$

Combining (15) and (18) gives

$$m(\mathbf{x}) = m^0(\mathbf{x}) + \int K(\mathbf{x}, \mathbf{x}) [\mathcal{N}(\mathbf{x}|b_1, c_1), \dots, \mathcal{N}(\mathbf{x}|b_M, c_M)] \beta (\mathbf{u} - \boldsymbol{\rho}) d\mathbf{x} \quad (19)$$

$$= m^0(\mathbf{x}) + \tilde{K}(\mathbf{x}, B) \beta (\mathbf{u} - \boldsymbol{\rho}) = m^0(\mathbf{x}) + \tilde{K}(\mathbf{x}, B) (\alpha - \beta \tilde{\mathbf{m}}_B^0) \quad (20)$$

where $\tilde{K}(\mathbf{x}, B)$ is defined in (14). Setting the prior mean function $m^0(\mathbf{x}) = 0$, we have $\tilde{\mathbf{m}}_B^0 = 0$. As a result, equation (12) holds.

From (18) it follows that

$$\frac{d^2 \log Z}{dm^0(\mathbf{x})dm^0(\mathbf{x}')} = -[\mathcal{N}(\mathbf{x}|b_1, c_1), \dots, \mathcal{N}(\mathbf{x}|b_M, c_M)] \beta \frac{d\boldsymbol{\rho}}{dm^0(\mathbf{x}')} \quad (21)$$

$$= -[\mathcal{N}(\mathbf{x}|b_1, c_1), \dots, \mathcal{N}(\mathbf{x}|b_M, c_M)] \beta [\mathcal{N}(\mathbf{x}'|b_1, c_1), \dots, \mathcal{N}(\mathbf{x}'|b_M, c_M)]^T \quad (22)$$

Based on the above equation and (16), we have

$$V(\mathbf{x}, \mathbf{x}') = K(\mathbf{x}, \mathbf{x}') - \tilde{K}(\mathbf{x}, B) \beta \tilde{K}(B, \mathbf{x}') \quad (23)$$

Thus (13) holds. \square

Blurring the Gaussian process by the local Gaussian distribution $\mathcal{N}(\mathbf{x}|b_i, c_i)$, we obtain the following corollary:

Corollary 2 *The projection of the blurred Gaussian posterior process $q(f)$ onto B is a Gaussian distribution with the following mean and covariance:*

$$\tilde{\mathbf{m}}_B = \hat{\mathbf{K}} \alpha \quad (24)$$

$$\tilde{\mathbf{V}}_B = \hat{\mathbf{K}} - \hat{\mathbf{K}} \beta \hat{\mathbf{K}}^T \quad (25)$$

where $\hat{K}_{ij} = \int \int \mathcal{N}(\mathbf{x}|b_i, c_i) K(\mathbf{x}, \mathbf{x}') \mathcal{N}(\mathbf{x}'|b_j, c_j) d\mathbf{x} d\mathbf{x}'$.

Now the remaining question is how to estimate $(\mathbf{u}, \boldsymbol{\Lambda})$, or equivalently (α, β) , such that the approximate posterior distribution on B well approximates the exact posterior. In the following section, we describe an inference method for the sparse model.

3.2 Inference by expectation propagation

Expectation propagation (Minka, 2001) has three steps, message deletion, data projection, and message updates, iteratively applied to each training point. In the message deletion step, we compute the partial belief $q^{\setminus i}(f; m^{\setminus i}, v^{\setminus i})$ by removing a message \tilde{t}_i (from the i -th point) from the approximate posterior $q^{\text{old}}(f|m, v)$. In the data projection step, we minimize the KL divergence between $\tilde{p}(f) \propto p(t_i; f) q(f; m^{\setminus i}, v^{\setminus i})$ and the new approximate posterior $q^{\text{old}}(f|m, v)$, such that the information from each data point is incorporated into the model. Finally, the message \tilde{t}_i is updated based on the new and old posteriors.

Based on (5), the sparse GP is an exponential family with features $(\mathbf{g}_B(f), \mathbf{g}_B(f)\mathbf{g}_B(f)^T)$. As a result, we can determine the sparse GP that minimizes $KL(\tilde{p}(f)|q(f))$ by matching the moments on $\mathbf{g}_B(f)$.

Similar to (15) and (16), the moment matching equations are

$$\tilde{\mathbf{m}}_B = \tilde{\mathbf{m}}_B^{\setminus i} + \tilde{V}^{\setminus i}(B, \mathbf{x}_i) \frac{d \log Z}{dm^{\setminus i}(\mathbf{x}_i)} \quad (26)$$

$$\tilde{\mathbf{V}}_B = \tilde{\mathbf{V}}_B^{\setminus i} + \tilde{V}^{\setminus i}(B, \mathbf{x}_i) \frac{d^2 \log Z}{(dm^{\setminus i}(\mathbf{x}_i))^2} \tilde{V}^{\setminus i}(\mathbf{x}_i, B) \quad (27)$$

$$\text{where } \tilde{V}^{\setminus i}(B, \mathbf{x})_j = \int \mathcal{N}(\mathbf{x}'|b_j, c_j) V^{\setminus i}(\mathbf{x}', \mathbf{x}) d\mathbf{x}' \quad (28)$$

Combining (24) and (26) gives

$$\alpha = \alpha^{\setminus i} + \mathbf{b} \frac{d \log Z}{dm^{\setminus i}(\mathbf{x}_i)} \quad (29)$$

$$\mathbf{b} = \hat{\mathbf{K}}^{-1} \tilde{V}^{\setminus i}(B, \mathbf{x}_i) = \mathbf{p}_i - \beta_{\setminus i} \tilde{K}(B, x_i) \quad (30)$$

$$\mathbf{p}_i = \hat{\mathbf{K}}^{-1} \tilde{K}(B, x_i) \quad (31)$$

where we use (13) to obtain the last equation in the second line.

From (25) we have $\beta = \hat{\mathbf{K}}^{-1}(\hat{\mathbf{K}} - \tilde{\mathbf{V}}_B)(\hat{\mathbf{K}}^{-1})^T$. Inserting (27) into this equation gives

$$\beta = \beta^{\setminus i} - \mathbf{b} \mathbf{b}^T \frac{d^2 \log Z}{(dm^{\setminus i}(\mathbf{x}_i))^2} \quad (32)$$

These equations define the projection update. This update can be equivalently interpreted as multiplying $q^{\setminus i}(f)$ by an approximate factor $\tilde{t}_i(f)$ defined as:

$$\tilde{t}_i(f) = \mathcal{N}(\sum_j p_{ij} \int f(\mathbf{x}) \mathcal{N}(\mathbf{x}|b_j, c_j) d\mathbf{x} | g_i, \tau_i^{-1}) \quad (33)$$

$$\tau_i^{-1} = (-\nabla_m^2 \log Z)^{-1} - \tilde{K}(\mathbf{x}_i, B) \mathbf{b} \quad (34)$$

$$g_i = m^{\setminus i}(\mathbf{x}_i) + (-\nabla_m^2 \log Z)^{-1} \nabla_m \log Z \quad (35)$$

The approximation factor $\tilde{t}_i(f)$ can be viewed as a message from the i -th data point to the sparse GP. To check the validity of this update, we compute

$$\tilde{Z} = \int \tilde{t}_i(f) q^{\setminus i}(f) df \propto \mathcal{N}(u_i | \mathbf{p}_i^T \tilde{\mathbf{m}}_B^{\setminus i}, \tau_i^{-1} + \mathbf{p}_i^T \tilde{\mathbf{V}}_B^{\setminus i} \mathbf{p}_i) \quad (36)$$

$$= \mathcal{N}(u_i | m^{\setminus i}(\mathbf{x}_i), \tau_i^{-1} + \tilde{K}(\mathbf{x}_i, B) \mathbf{b}) \quad (37)$$

which has the same derivatives as the original $Z = \int t_i(f) q^{\setminus i}(f) df$. Therefore, the multiplication $\tilde{t}_i(f) q^{\setminus i}(f)$ leads to the same $q(f|m, v)$. In other words, we have $\tilde{t}_i(f) \propto q(f)/q^{\setminus i}(f)$.

To delete this message, we multiple its reciprocal with the current $q(f)$. Using the same trick as before, we can solve the multiplication using the following moment matching equations:

$$\mathbf{h} = \mathbf{p}_i - \beta \tilde{K}(B, \mathbf{x}_i) \quad (38)$$

$$\tilde{Z}_d = \int \frac{1}{\tilde{t}_i(f)} q^{\setminus i}(f) df \propto \mathcal{N}(u_i | \mathbf{p}_i^T \tilde{\mathbf{m}}_B^{\setminus i}, -\tau_i^{-1} + \tilde{K}(\mathbf{x}_i, B) \mathbf{h}) \quad (39)$$

$$\frac{d^2 \log \tilde{Z}_d}{(dm(\mathbf{x}_i))^2} = -(-\tau_i^{-1} + \tilde{K}(\mathbf{x}_i, B) \mathbf{h})^{-1} \quad (40)$$

$$\frac{d \log \tilde{Z}_d}{dm(\mathbf{x}_i)} = (-\frac{d^2 \log Z}{(dm(\mathbf{x}_i))^2})(g_i - \tilde{K}(\mathbf{x}_i, B) \alpha) \quad (41)$$

$$\alpha^{\setminus i} = \alpha + \mathbf{h} \frac{d \log \tilde{Z}_d}{dm(\mathbf{x}_i)} \quad (42)$$

$$\beta^{\setminus i} = \beta - \mathbf{h} \mathbf{h}^T \frac{d^2 \log \tilde{Z}_d}{(dm(\mathbf{x}_i))^2} \quad (43)$$

The inference algorithm is summarized in Algorithm 1.

-
1. Initialize $q(f)$, g_i , and τ_i all to be 0.
 2. Loop until the change over all g_i , and τ_i is smaller than a threshold

Loop over all training data point \mathbf{x}_i

Deletion. Compute $\alpha^{\setminus i}$ and $\beta^{\setminus i}$ for $q^{\setminus i}(f)$ via (42)(43)
Projection. Compute α and β for $\mathbf{q}(f)$ via (29)(32)
Inclusion. Update g_i , and τ_i for the message \tilde{t}_i via (34)(35)
-

Algorithm 1: Expectation Propagation Inference for Sparse GP

3.3 Regression

Given the linear regression likelihood (2), the quantities in the projection step (29)(32) are

$$\frac{d \log Z}{dm(\mathbf{x}_i)} = \frac{y_i - m^{\setminus i}(x_i)}{v_y + v^{\setminus i}(x_i, x_i)} \quad \frac{d^2 \log Z}{(dm(\mathbf{x}_i))^2} = \frac{-1}{v_y + v^{\setminus i}(x_i, x_i)} \quad (44)$$

3.4 Classification

Given the classification likelihood (3) where $\phi(\cdot)$ is the step function, the quantities in the projection step (29)(32) are

$$z = \frac{m^{\setminus i}(x_i)y_i}{\sqrt{v^{\setminus i}(x_i, x_i)}} \quad (45)$$

$$Z = \epsilon + (1 - 2\epsilon)\psi(z) \quad (46)$$

$$\frac{d \log Z}{dm(\mathbf{x}_i)} = \gamma y_i \quad (47)$$

$$\frac{d^2 \log Z}{(dm(\mathbf{x}_i))^2} = \frac{\gamma(m^{\setminus i}(x_i)y_i + v^{\setminus i}(x_i, x_i)\gamma)}{v^{\setminus i}(x_i, x_i)} \quad (48)$$

where $\gamma = \frac{(1-2\epsilon)\mathcal{N}(z|0,1)}{Z\sqrt{v^{\setminus i}(\mathbf{x}_i)}}$ and $\psi(\cdot)$ is the standard Gaussian cumulative distribution function.

4 Related work

Most work in sparse GP approximation has been for regression with Gaussian noise, where the exact posterior distribution is a GP whose moments are available in closed form. The goal in that case is to approximate the exact GP with a simpler GP. Quinonero-Candela and Rasmussen (2005) compared several such approximations within a common framework, highlighting the FITC approximation as an improvement over DTC. Walder et al. (2008) later generalized FITC to include basis-dependent lengthscales, and showed that this approximation fits into the framework of Quinonero-Candela and Rasmussen (2005).

For classification with GPs, few papers have imposed an explicit sparsity constraint on the GP posterior. One reason for this is that sparse GP approximations have only been understood from a regression perspective. Csato and Opper (2000) showed how online GP classification could be made sparse, using an approximation equivalent to FITC. However, the FITC approximation was introduced as a subroutine, not presented as a part of EP itself. Csato (2002) later gave a batch EP algorithm, where a different sparse GP approximation (DTC) was used as a subroutine within EP. The software distributed by Csato can run either online or batch and has an option to use either FITC or DTC. In this paper, we clarify Csato and Opper’s work by showing that the FITC approximation can in fact be viewed as part of the overall EP approximation. Furthermore we extend it to include VSGP approximation.

5 Experiments

We evaluate the new sparse GP method on both synthetic and real world data for regression and classification tasks and compare its predictive performance with alternative sparse GP methods. We use the RBF kernels for all the experiments. To all the sparse GP models, we need to choose basis centers. Instead of using the evidence framework, we use the entropy reduction criteria (equations are omitted to save space) to search over the hierarchical partitions of the data. We then use the cluster centers as the basis points and use the data partitions to compute the local covariance matrix. Our focus is on the approximate inference method so we use the same basis points for all the methods being evaluated.

We first examine the performance of the new method on synthetic data for the regression task. Since we can control the generative process of the synthetic data, it is easier for us to gain insight into how the method performs. For regression, we sample 100 data points along a circle with some additive Gaussian noise. The output in different quadrant has different values plus certain additive Gaussian noise.




Figure 1: Illustration on simple circle data. The heatmaps represent the values of the posterior means of different methods. Red ellipses and crosses are the mean and the standard deviation of local covariances for the sparse GP model. The white dots are the training data points. EP estimation uses the full local covariance matrix in (d), significantly improving the estimation accuracy along the circle.

The mean of the exact posterior distribution of the full GP is shown in figure 1.a. The posterior means of FITC, SOGP, and the new method with sphere and full local covariance matrices. Four basis points are chosen as the centers learned by K-means clustering. Note that when the local covariance matrices become sphere matrices, our method reduces to the multiscale method (Walder et al., 2008). For FITC, we performed cross-validation to choose the right kernel width for the RBF kernel. As shown in the figure, the use of the full local covariance matrix improves the estimation accuracy.

Then we perform test FITC, SOGP, and the new methods with sphere and full local covariance matrices on a subset of the pumadyn-32nm dataset. On this dataset, however, we do not see the

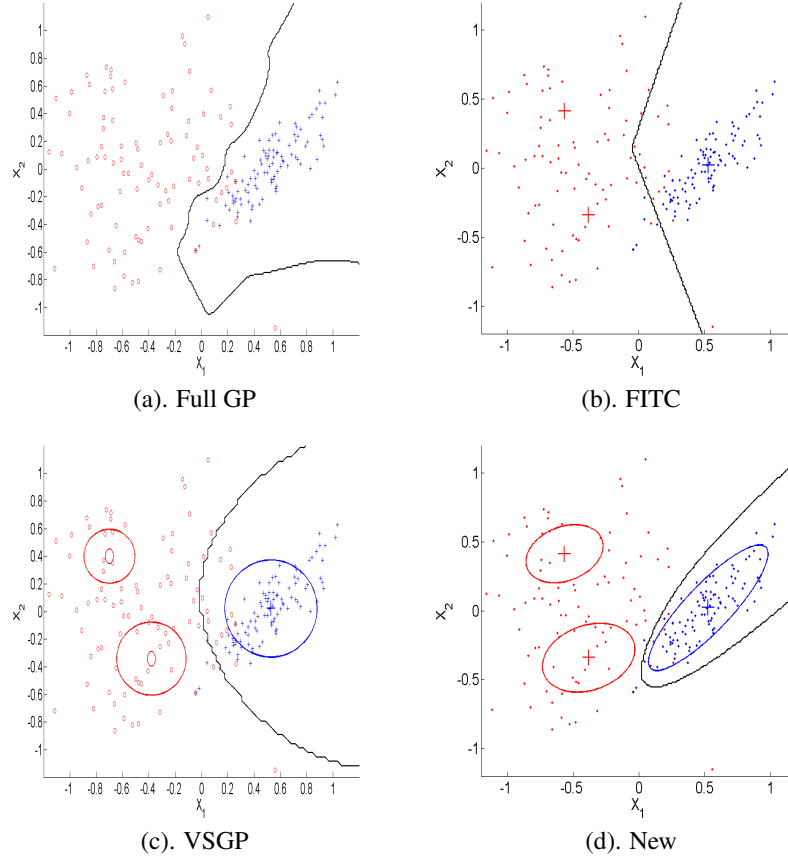


Figure 2: Classification on synthetic data. The blue and red ellipses show the standard deviation of local covariances for the sparse GP model. The black curve is the decision boundary. With only three basis points, the true, complex decision boundary in (a) is well approximated by an ellipse by our method (d).

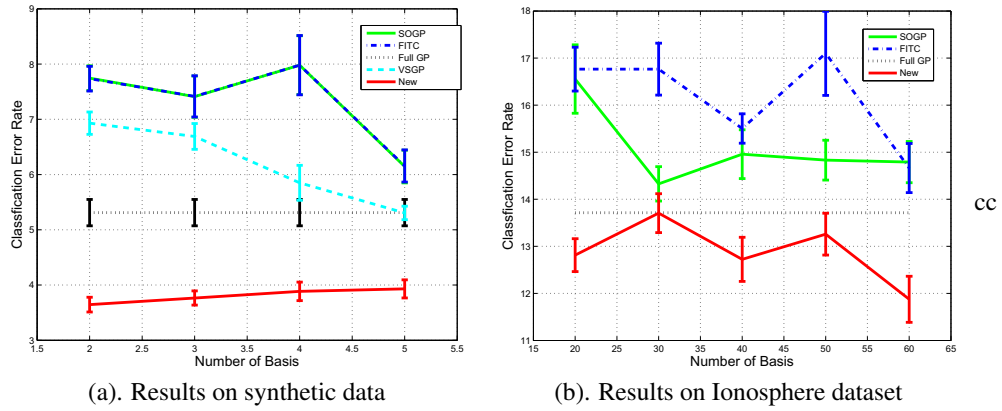


Figure 3: Classification on synthetic and real data.

further improvement of our new method over FITC. This is consistent with the results reported in (Walder et al., 2008).

We then test all the methods on synthetic data for a classification task. Each class of the data is sampled from a Gaussian and shown in Figure 3. It also shows the decision boundaries and the basis points. Figure 3 illustrates the advantages of our new approach. Since we simply use K-means to choose the basis points, further improvement for FITC is possible. We then sample 200 points for training and 2000 for testing. The experiments are repeated 10 times and the results are shown in Figure 3. In this figure, we also report results of SOGP.

Finally we tested our methods on a standard UCI dataset, Ionosphere. The results are summarized in Figure 3. Again, we outperform the previous methods.

6 Conclusions

In this paper, we present a new perspective to understand sparse GP models using the expectation propagation framework and develop new inference methods based on this perspective. Empirical results demonstrate improved prediction accuracy with the new extensions.

References

- Chu, W., Sindwhani, S., Ghahramani, Z., & Keerthi, S. S. (2006). Relational learning with gaussian processes. *Advances in Neural Information Processing Systems 18*.
- Csato, L. (2002). *Gaussian processes - iterative sparse approximations*. Doctoral dissertation, Aston University.
- Csato, L., & Opper, M. (2000). Sparse representation for Gaussian process models. *Advances in Neural Information Processing Systems 13* (pp. 444–450). MIT Press.
- Deisenroth, M. P., Rasmussen, C. E., & Peters, J. (2009). Gaussian process dynamic programming. *Neurocomputing*, 72, 1508–1524.
- Lawrence, N., Seeger, M., & Herbrich, R. (2002). Fast sparse Gaussian process methods: The informative vector machine. *Advances in Neural Information Processing Systems 15* (pp. 609–616). MIT Press.
- Minka, T. P. (2001). Expectation propagation for approximate Bayesian inference. *In Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence* (pp. 362–369).
- Quinonero-Candela, J., & Rasmussen, C. E. (2005). A unifying view of sparse approximate Gaussian process regression. *Journal of Machine Learning Research*, 6, 1935–1959.
- Rasmussen, C. E., & Williams, C. K. I. (2006). *Gaussian processes for machine learning*. The MIT Press.
- Snelson, E., & Ghahramani, Z. (2006). Sparse gaussian processes using pseudo-inputs. *Advances in Neural Information Processing Systems 18* (pp. 1257–1264). MIT press.
- Teh, Y. W., Seeger, M., & Jordan, M. I. (2005). Semiparametric latent factor models. *The 8th Conference on Artificial Intelligence and Statistics (AISTATS)*.
- Walder, C., Kim, K. I., & Schölkopf, B. (2008). Sparse multiscale Gaussian process regression. *ICML '08: Proceedings of the 25th international conference on Machine learning* (pp. 1112–1119). New York, NY, USA: ACM.